PERFORMANCE ASSESSMENT OF VISION BASED HAZARD AVOIDANCE DURING LUNAR AND MARTIAN LANDING

B. Parreira(1), J. F. Vasconcelos (1), R. Oliveira(1) , A. Caramagno(1) , P. Motrena(2) , J. Dinis(2) , J. Rebordão(2)

(1)DEIMOS Engenharia, Av. D. João II, Lote 1.17.01-10º, Lisboa, Portugal, Email: baltasar.parreira@deimos.com.pt
(2) Faculty of Sciences of the University of Lisbon – Laboratory of Optics, Lasers and Systems, Estrada do Paço do Lumiar 22 Building D, Lisboa, Portugal, Email: pjmpereira@fc.ul.pt

ABSTRACT

Hazard Avoidance is a key technology for a safe landing of future planetary missions. During Hazard Avoidance, sensors and computers onboard the lander are used to detect hazards in the landing zone, autonomously select the most suitable region for landing, and generate the trajectory that retargets the lander to the safer landing site. In this paper, vision-based hazard avoidance algorithms are briefly described, and the results of testing under two realistic simulated scenarios representing a landing on Mars and Moon are presented. Results show that the developed Hazard Avoidance algorithms are effective at detecting hazards and guiding the lander to a safe landing site.

1. INTRODUCTION

Technologies for planetary landing have been studied and developed since the late fifties during the Moon race, which culminated in 1969 with the first human landing on the Moon. Nowadays, instead of humans, small probes/landers are sent to distant planetary bodies. These landings are often performed by a pre-programmed time sequence of events that bring the lander to a full stop in a desired area at the planet surface (e.g. 'pathfinder-type' is an open-loop landing with airbags; 'Viking-type' a semi-automatic landing).

Future exploration missions envisage landing on planetary surfaces that are not well known apriori, or in areas that are not as flat and hazard free as the nominal selected Landing Sites (LS) of past exploration missions. Landers also tend to become smaller and lighter, not so robust to surface hazards. Autonomous pinpoint soft-landing systems that include Hazard Avoidance (HA) capability are therefore required to guarantee safe landing.

A HA system is responsible for the detection of any hazards that put in risk the landing mission, and path-planning to avoid the detected hazards. Hazard detection implies the lander to be equipped with proper sensing devices. In the frame of this study, an optical sensor (onboard camera) is used to detect hazards (e.g. shadows, boulders, high slopes) in the landing zone.

This paper briefly describes the developed HA algorithms (already described in more detail in [2]), focusing on the functional consolidation aspects aimed to increase the system's technology readiness level, and presents the work done in the testing of the algorithms on a realistic landing on Moon and Mars.

2. HAZARD AVOIDANCE

2.1 Hazard Avoidance Concept

Hazard Avoidance algorithms are responsible for the detection and avoidance of hazards by retargeting, when necessary, to a 'new' LS on a hazard-free area. The choice of the 'new' LS is such that it can be reached with the onboard fuel while complying with the vehicle and mission requirements and constraints. The HA algorithms encompass the following functions:

- **Hazard Mapping**: refers to the process of analysing terrain topography and detecting hazards through image processing algorithms applied to the monocular optical images taken by the onboard camera.
- **Piloting**: refers to the concepts of data fusing, planning and decision-making used for the selection of a safe LS.
- **Guidance**: refers to the concepts used to steer the spacecraft till the LS.

Fig. 1 presents a sketch of the HA functions within the complete GNC scheme, which are described in more detail in the next subsections.

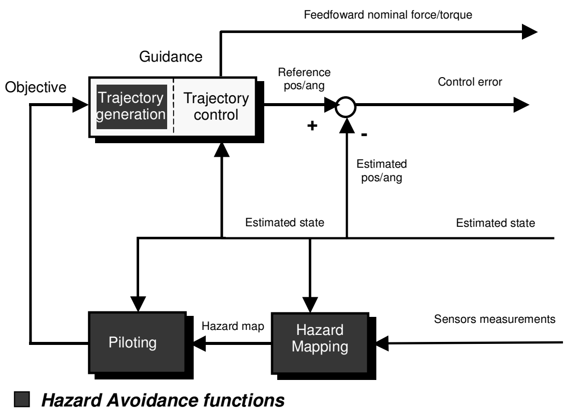

Fig. 1. Hazard avoidance functions

2.2 Hazard Mapping

Hazard mapping (HM) is the process of assigning a hazard score to each pixel of the image, which measures in both absolute and relative terms how risky it is for landing. Hazards are computed within a region of interest built by Piloting around a candidate landing site.

Hazard map components are slopes, shadows, and texture. Due to mission constrains (and study specifications) hazard components must be extracted exclusively from panchromatic monocular images. Fig. 2 shows a typical image as provided by PANGU image generator tool.

Fig. 2. Terrain image, generated with PANGU

Shadows can be extracted by automatic thresholding using the image grey level histogram and are graded linearly between the minimum grey level and such threshold.

Texture is assessed using the grey level standard deviation within a square patch.

Slope is estimated by first building a digital elevation model using shape from shading (SFS) techniques, supplemented by spatial filtering to remove striping effects and to reintroduce coherence between consecutive lines.

Slopes and shadows are absolute hazard indicators, although texture is normalized to the maximum value within the region of interest.

Fig. 3 presents (top to bottom) the shadow, texture and slope maps as have been extracted from the camera image reported in Fig. 2.

The scores of each of these factors are then combined into a single score that measures the hazard of landing for each pixel. This resulting hazard map is scaled between 0 (no risk) and 1 (maximum risk) and is used by Piloting to derive sufficient conditions for retargeting.

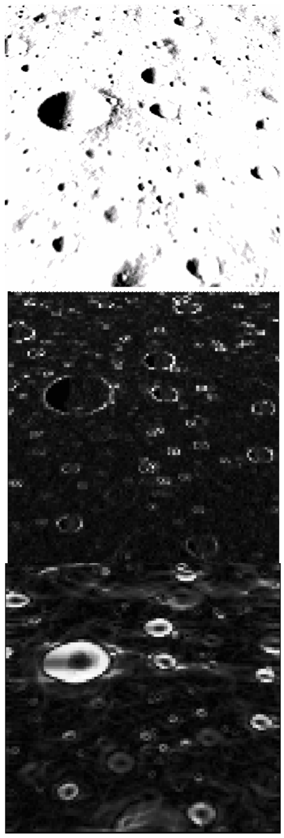

Fig. 3. Shadows, texture and slope maps

2.3 Piloting

Piloting is the function of the HA system that is responsible for evaluating if the current landing site is considered unsafe and, whenever this happens, to provide a new safe LS. Piloting selects a LS that is, not only ‘safer’ but also reachable. This attainability is addressed considering system requirements and mission constraints such as: available on-board propellant, LS visibility, and avoiding any crashing condition.

The Piloting function follows three main phases:

1. **Data collecting**: all the relevant information is gathered and organized in a particular 'map' format. Information regarding the hazards on the landing zone, the fuel consumption till a candidate LS, the required thrust profile to fly the retargeting trajectory, the candidate LS visibility along the trajectory, etc. is collected and used to define the attainable retargeting area, which is the region on the planet's surface that is suitable for retargeting.
2. **Data fusing**: the gathered information is combined/merged into a Global Score Map.
3. **Decision-making**: given the Global Score Map a decision is made on the 'safest' LS, which is then provided to the Guidance module.

Data collecting

During this algorithm phase, relevant information is gathered and organized in matrices, whose elements correspond to the pixels of the image taken by the camera. The elements represent risk/cost coefficients with scores between zero and one (0 = low risk/cost, 1 = high risk/cost).

Firstly, the algorithm computes the Attainable Retargeting Area Map. This map restricts the landing zone to a smaller 'retargeting area', which corresponds to the region, on the landing zone, that can be reached fulfilling mission requirements while complying with constraints.
In order to define this region, the coordinates on the planet surface corresponding to each image pixel are firstly computed. Then, the distance between each pixel and the current LS is determined. Finally, a score is assigned giving preference to landing sites closer to the nominal one (due to mission design and/or for scientific reasons).
For those pixels within the required retargeting area (that means Distance Score lower than 1) the lander trajectory to reach the corresponding LS is estimated using a simplified guidance scheme. The estimation of the trajectory allows evaluating if the trajectory is 'flyable', in a sense that it meets the lander manoeuvrability constraints. These encompass fuel consumption, trajectory acceleration compatibility with propulsion system capabilities, and LS visibility.

One key aspect of the system maturation concerns the real-time operation constraints. Effectively, in view of the real-time implementation of the HA system, it is necessary to consider the time that each of the HA functions will take to execute, and the impact this will have at the functional level of the system.
In particular, this will affect the calculation of the attainable region, as this is done for the time instant when the processing is initiated, not for the time instant when the actual retargeting can be commanded.
To avoid discrepancy in the attainable region, Piloting propagates the spacecraft state until the time of retargeting, and computes the attainable region for that spacecraft state. The propagation is done by projecting ahead in time (using the simplified guidance scheme) the reference trajectory to the nominal LS.

The Attainable Area Map deals with the reachability aspects of the landing sites available in the terrain. For the site safety aspects, Piloting computes the Risk Map, which essentially constitutes its interpretation of the information provided by the Hazard Mapping function.
Two different approaches for Risk Map computation have been implemented:

- **Smoothed Risk Map:** The risk score of a given LS reflects the HM score of itself and also the scores of the surrounding sites inside an area defined by the GNC dispersion. This approach is used in the initial phase of the landing to select a landing area that is globally safe.
- **Distance-based Risk Map:** The candidate LS score reflects its distance to the nearest hazardous LS. This will actively avoid ground obstacles by maximizing ground distance between lander and nearest hazard. (i.e. it will promote the center of the hazard-free areas). This approach is used in the final phase of landing in order to avoid smaller terrain hazards (e.g. boulders).

Fig. 4 illustrates the Risk Maps computed with both approaches (lighter colour indicate hazardous areas).

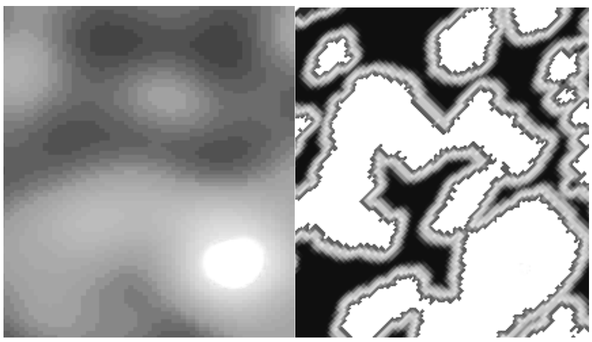

Fig. 4. Risk Map computation: smoothed (left), distance-based (right).

Data Fusing

This function combines all the gathered information into a global score map,

GlobalMap = *f*(AttainableRegionMap, RiskMap)

where *f* is the data fusing function, defined such that the global score is "0" for "bad" LS candidates and "1" for "good" LS candidates. It is only necessary for one score to be bad "0" to lead the Global Score to a bad value.

Decision Making

The Decision Making function chooses which LS is to be provided to the guidance function. It is clear that the best LS corresponds to the one having the highest global score. However, retargeting only occurs if:

- the global score of a candidate site is significantly better than the one of the current target;
- the same candidate site is chosen for some iterations.

2.4 Guidance

Within a HA system it is the task of the guidance function to steer the spacecraft from the current state to a specific target state, which can change during flight (through retargeting).
The guidance algorithm that has been developed is a terminal point guidance, based on the E-Guidance [1]. With the assumptions of a flat surface and constant gravity field, the method computes the guidance acceleration command by solving, for each axis component, a two-point boundary value problem to guide the vehicle from its current state to the desired target state, in a specific time-to-go (T_{go}).
T_{go} is the synchronizing variable that ensures simultaneous solutions of all dimensions of the total guidance problem. T_{go} is computed to erode the vehicle's velocity with a constant acceleration.
Algorithm follows with the computation of the so-called E-matrix and the polynomial coefficients c_1 and c_2 for the acceleration command:

$$\begin{bmatrix} c_1 \\ c_2 \end{bmatrix} = \begin{bmatrix} \frac{4}{T_{go}} & -\frac{6}{T_{go}^2} \\ -\frac{6}{T_{go}^2} & \frac{12}{T_{go}^3} \end{bmatrix} \begin{bmatrix} V_{Y_{LS}} - V_Y \\ Y_{LS} - \left(Y + V_Y \cdot T_{go}\right) \end{bmatrix} \quad (1)$$

Note that this is performed for each axis (such that the boundary conditions are satisfied in all axes) and therefore we obtain the polynomial coefficients vectors: $\vec{\mathbf{c}}_1$ and $\vec{\mathbf{c}}_2$. Also note that these coefficients are computed at each guidance step, up to the moment where they are frozen to their previous values (whenever $T_{go} < T_{go\ Freeze}$) in order to avoid numerical problems with E-matrix computation.
Finally the thrust acceleration command vector is computed with

$$\vec{\mathbf{a}}_T(t) = \vec{\mathbf{c}}_1 + \vec{\mathbf{c}}_2\left(T_{go} - t\right) + \vec{\mathbf{g}} \quad (2)$$

where $\vec{\mathbf{g}}$ is the planet gravity acceleration. The required thrust force command is computed as:

$$\vec{\mathbf{T}} = \left(m_{dry} + m_{fuel}\right)\vec{\mathbf{a}}_T \quad (3)$$

where m_{dry} and m_{fuel} are the lander dry mass and fuel mass, respectively.

3. PERFORMANCE ASSESSMENT

3.1 Overview

Two mission scenarios are considered: a Mars scenario based on Mars Sample Return Mission, and a Moon scenario based on the Lunar Lander Mission. Landing is expected to take place in demanding conditions for the vision-based Hazard Avoidance system. In particular, for the Moon scenario, a very low Sun elevation angle (1.5º) is foreseen.
For the purpose of developing and testing the HA algorithms, the current study focus on a particular segment of the landing phase, which starts at the High Gate (HG), with the first acquisition of images of the nominal landing area; and ends at the Low Gate (LG), point where the LS visibility conditions are no longer satisfied.
The developed algorithms have been integrated in the 'Vision-Based Navigation Analysis Tool' (VBNAT) (developed by EADS Astrium, and provided by ESA) for six degrees of freedom, high fidelity, closed-loop testing in a planetary landing scenario. The VBNAT simulator includes an interface with a synthetic terrain generator – PANGU (developed by University of Dundee) which is used for image generation.
In order to validate correct operation, the HA algorithms have been tested in each scenario using the VBNAT simulator. The test cases consist in simulating the landing mission, targeting a nominal LS that is unsafe, and checking if a retargeting is commanded towards a safe LS. The results obtained so far are with ideal navigation. Nevertheless, these results still provide a meaningful indication of the HA system's capability to select a safe LS, with the navigation error being expected to have a more important impact on the

capability of the GNC system to guide the lander until the designated LS.
The following subsections present the terrain characteristics being considered for the test campaign and the results achieved in each scenario.

3.2 Terrain characteristics

The terrain used to test HA operation is a challenging one, not only because of the shallow illumination that causes large shadows to be cast throughout the terrain, but also because of the very high number of dangerous boulders. Fig. 5 presents a view of the area surrounding the nominal LS, where it can be visible a large central crater 80m in diameter, and several large boulders spread around the terrain.

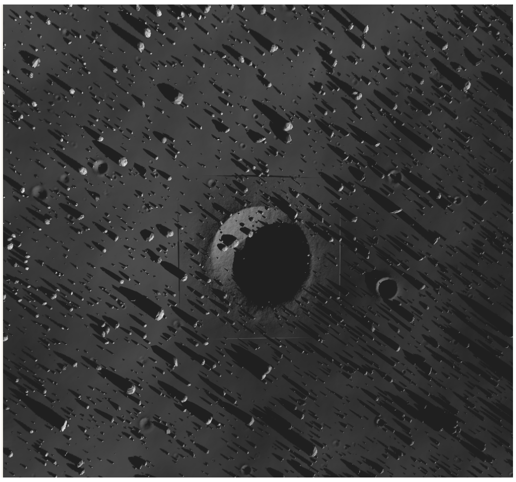

Fig. 5. Detailed view of nominal LS and surrounding area.

In order to check if a landing simulation has been successful, it is necessary to know if the final location reached by the lander is safe or not. To this end, a specific tool has been developed that allows determining the Reference Hazard Map (RHM) of the terrain. The RHM contains the real safety level of each LS available in the terrain and is obtained by processing the PANGU model data. Sites that are hazardous due to shadows are identified by generating a highly contrasted image of the terrain. The locations that contain hazardous boulders are identified by processing a PANGU data file that contains a list of all the boulders in the terrain. To determine the areas with dangerous slopes, the terrain digital elevation model (DEM) is first extracted from PANGU, and then local slopes are computed.
Fig. 6 presents the RHM for the Moon scenario with the corresponding camera image. The white regions in the RHM correspond to hazardous sites which constitute about 45% of all sites available. The streaking pattern that can be observed is caused by the long shadows cast by large boulders, which in turn are due to the very shallow illumination.

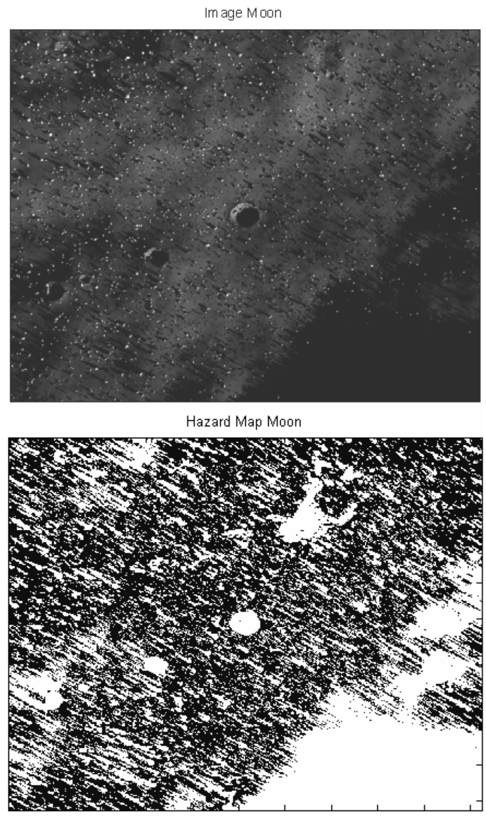

Fig. 6. Reference Hazard Map (bottom), and corresponding camera image (top).

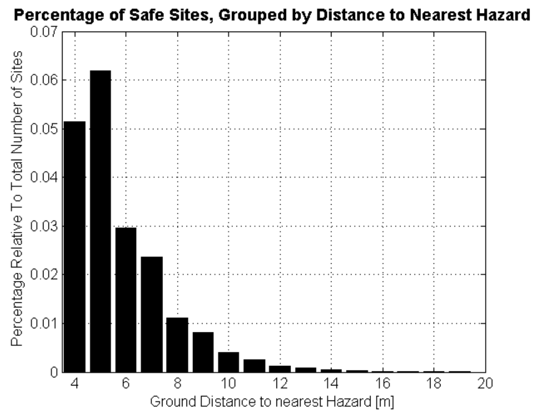

Fig. 7. Histogram of the distance between each safe landing site and the nearest hazard in Moon scenario.

Fig. 7 presents a histogram of the distance between each safe landing site and the nearest hazard. Note that only sites with distances larger than 4m are plotted since this is the minimum clearance in order to assure that the lander footprint does not overlap a hazardous location. This histogram effectively illustrates that only very few safe locations on the terrain have a reasonable amount of clearance to hazards. Sites with a clearance above 15m are almost non existent.

3.3 Mars scenario results

The HA system operation in the Mars scenario was simulated using VBNAT. The nominal LS targeted at the start of landing is unsafe, but the system commanded 2 retargetings which brought the lander to a safe location 190m away from the initial LS. The first retargeting was commanded at the early phase of landing and changed the targeted LS to a distant area that is globally safe, Fig. 8. The second retargeting occurred later and was of smaller amplitude.

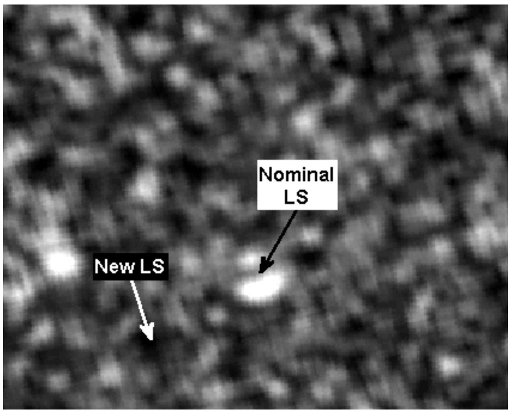

Fig. 8. Hazard Map corresponding to first retargeting in Mars scenario.

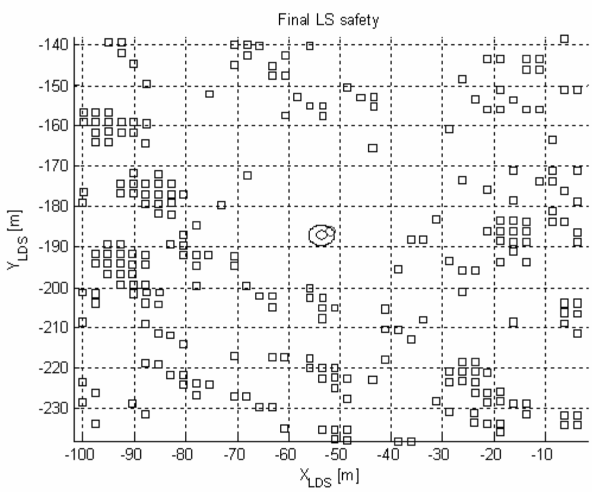

Fig. 9. Safety of final reached position (Mars scenario). Circles represent the final selected/reached positions; squares mark hazardous locations.

Fig. 9 presents the RHM of the terrain, together with the final selected LS location and also the final position reached by the lander. It can be seen that the system targets the centre of a relatively large clearing, with the finally selected/reached positions being more than 13m away from the nearest hazard.

3.4 Moon scenario results

In the Moon landing scenario the system also commands 2 retargetings that divert the lander to a safe position 80m away. Similarly to Mars, the first retargeting occurs early and is effective in selecting a globally safe area. The second retargeting is of smaller amplitude and is motivated by the need to avoid a dangerous boulder that only becomes visible later in the descent. Fig. 10 presents the HM corresponding to this retargeting, where it is possible to see the isolated boulder which was originally targeted, and the new selected LS which is towards the centre of a nearby safe area. This second retargeting illustrates well the benefit of introducing the distance-based risk map for the final phase of landing.

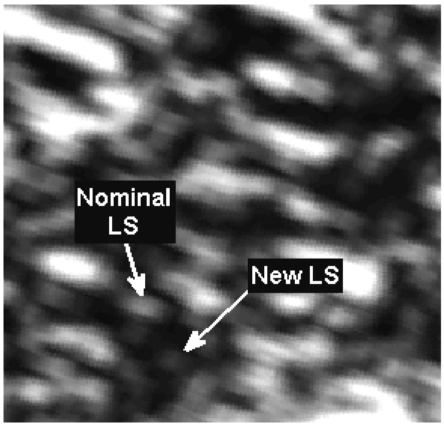

Fig. 10. Hazard Map corresponding to second retargeting in Moon scenario.

The safety level of the final reached position is presented in Fig. 11. The finally targeted position is within a relatively large clearing, with the nearest hazard over 11m away. When comparing with the terrain distance histogram presented in Fig. 7, it can be seen that the LS indicated by the system is within the 1% best sites available in the terrain.

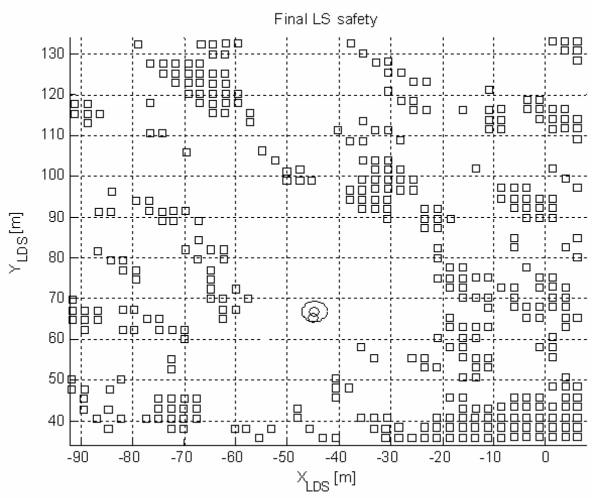

Fig. 11. Safety of final reached position (Moon scenario). Circles represent the final selected/reached positions; squares mark hazardous locations.

4. CONCLUSIONS

The presented HA system is currently being developed in order to consolidate the design and increase its maturity, with the ultimate objective of having a definitive implementation in a landing mission.

The design progresses presented in this paper follow this development path. The spacecraft state propagation implemented in Piloting aims to overcome algorithmic limitations that become evident when considering real-time operation of the system. The introduction of a new distance-based risk map improves the system performance in very hazardous terrains.

Developing tools that allow the correct validation of the HA system performance is also of high importance, with the RHM computation tool being a key asset.

The presented test cases demonstrate that the system is capable of selecting a landing site that is, not only safe, but also among the best available for landing.

5. ACKNOWLEDGEMENTS

The work presented in this paper was done under ESA/ESTEC contracts for the study of Visual Based Relative Navigation Techniques Framework (VBRNAV) (ESA/ESTEC contract No. 18038/04/NL/JA), and also of Vision Based Hazard Avoidance System Experiment (HASE).

6. REFERENCES

1. G.W. Cherry, “A general, explicit, optimising guidance law for rocket-propelled spaceflight”, AIAA paper no. 64–638, 1964.
2. F. Câmara et al., “Design and Performance Assessment of Hazard Avoidance Techniques for Vision Based Landing”, AIAA-2006-6593, AIAA Guidance, Navigation, and Control Conference and Exhibit, Keystone, Colorado, Aug. 21-24, 2006.
3. B. Parreira et al, “Consolidated Performance Assessment of Hazard Avoidance Techniques for Vision Based Landing”, AIAA-2007-6854, AIAA Guidance, Navigation, and Control Conference and Exhibit, Hilton Head, South Carolina, Aug. 20-23, 2007.
4. B. Parreira et al., “Hazard Avoidance for Planetary Landing: GNC Design and Performance Assessment”, GNC 2008 – 7th International ESA Conference on Guidance, Navigation & Control Systems, Tralee, County Kerry, Ireland, June 2008